SSD review report (3/14/08)

Committee: R. Bellwied (chair), D. Cebra, J. Engelage, J. Landgraf, G. Visser

The proposal to upgrade the readout of the SSD in STAR has been reviewed. In general our committee came to the conclusion that the project is essential for the success of the physics program of the joint Silicon upgrades (HFT, FGT, SSD) in STAR and therefore needs to be supported. In certain areas this initial proposal was too sketchy to be properly reviewed and we will point out the areas where follow-up reviews will be necessary.

We assume that this project is sufficiently inexpensive that it will be funded through sources which do not require the standard CD process as set up by the Department of Energy. Some of our recommendations will therefore reflect the necessity to keep a certain oversight on the project's progress through additional reviews. In the following we will address the specific charges to the committee by the STAR management.

Issues regarding multiple re-installation plan

The committee agrees with the SSD group's decision to not re-install the SSD in the summer of 2008. Therefore the first re-installation of the detector can happen during the 2009 shutdown. The committee recommends an engineering run after the 2009 shutdown, using three ladders with the new readout configuration, and then install the remaining 17 new ladders during the following shutdown. This two year installation schedule could move back by one year due to funding constraints, but our main recommendation is to not use any old readout components in any future SSD run.

We believe the main physics argument is still the role the SSD will play in an integrated tracker after the HFT and IST have been installed. The only other physics driver could be usage of the SSD for displaced vertices in Run 10, which will presumably be the first low-energy heavy ion run. In the present scheme this run would serve as an engineering run and the first real physics run with an upgraded SSD would be Run 11, which is likely a pp run. The group stated that the SSD could help the FGT physics program in Run 11, and we encourage STAR to find funding for the SSD and FGT projects so that both projects can be completed in time for Run 11. The first heavy ion physics results of a fully upgraded SSD will likely be produced in Run 12, which potentially is the second low energy physics run at RHIC.

Technical specifications

The physics-driven requirements for this detector are well understood, because the SSD can not be used as a stand-alone detector and therefore its requirements are driven by a.) the HFT requirements and b.) the DAQ requirements. There is no attempt to replace the Silicon wafers on the SSD, therefore the resolution capabilities are fixed and have been folded into the HFT performance estimates. This proposal exclusively deals with an

upgrade to the electronics, which also requires an upgrade to the cooling system and the mechanical structure in order to be integrated into the HFT/FGT support system. The electronics upgrade is driven by the increased DAQ requirements for the higher luminosity running of STAR in the next decade.
In general we believe that the proposed mechanical and electronics implementations will meet those requirements, although in parts of the proposal the necessary level of detail is missing and needs to be provided by the group and then reviewed by a technical committee by the end of this year.
We also believe that the proposed scope of the upgrade is justified by the physics driven requirements, although again there is a certain level of detail missing in the funding estimate, mostly because the actual new readout chain has not yet been fully planned and the work breakdown schedule has not been fully resource loaded. The manpower loading is sufficiently detailed and realistic, but the component and material loading has not been completed at this stage.

In the following we list a series of issues that need to be followed up by a technical review committee later this year or as soon as the final design is at hand.

a.) The new readout chain has not been fully defined. In particular, although the general ADC requirements are known, the sample rate and bit resolution are not stated in the proposal, and no candidate commercially available ADC's are named. The proposal initially gave the impression that the ADC and RDO electronics would be adopted from ALICE, but the only component that is actually an ALICE component is the DDL link which has been developed by the Budapest group in ALICE. The uncertainty in the readout chain design leads to an uncertainty in the power consumption budget (the proposal only states that the power budget will be likely higher than the present SSD budget), which in turn leads to an uncertainty in the layout of the detector cooling system, which leads to an uncertainty in the integration and the radiation length budget of the detector and the support system. The review committee anticipates that the increase in the power budget (which is presently around 400 W, i.e. close to the air system maximum for the on-detector electronics) will be moderate and that the SSD group will be able to cool the detector with a modified air cooling system rather than water or evaporative cooling. In fact, we would expect that recent developments in pipeline ADC's and the necessary support circuits for the commercial market may allow even to decrease the power of the ADC boards relative to the old system. But the power dissipation needs to be confirmed and reviewed in due time, because an alternate cooling system would significantly change the detector layout, the radiation length budget and the cost of the upgrade.

b.) If air cooling can be used then the modification to the existing cooling system will simply require more stable hoses and connectors. The biggest problem in the operation of the existing SSD was due to problems with kinked air hoses. Still, the new hose material and design for the air supplies has to be defined and reviewed, in particular with respect to radiation length constraints on the support cone, i.e. in forward direction.

c.) Although there is no reason to expect the particular L2 abort problems which plagued the existing system in a new design, the group needs to re-evaluate the L2 abort

functionality after the final design of the RDO chain is complete. The upgraded SSD readout system is required to conform to STAR trigger/DAQ requirements, including handling abort commands – we stress that this requirement presently remains in place even if L2 might be phased out in the DAQ1000 era. Additionally, while it is permissible for the RDO's to drop events in the rare case that the trigger/SSD interface is violated, the design should ensure that the RDO's do not ever become unresponsive or deliver mal-formed packets to DAQ. The committee also recommends that, if practical, a debugging facility in the RDO's to be able to write out a history of received trigger commands and tokens would be tremendously useful.

d.) The final readout speed can only be evaluated after a detailed design has been completed. The follow-up technical committee needs to review more detailed diagrams of the RDO chain from wafer to connection board to ADC board to RDO board to DAQ. The raw data volume (~675 kB) in conjunction with the anticipated low occupancy estimates could theoretically be managed with zero-suppression alone but the committee strongly recommends also buffering the data on the RDO boards. This will allow for the most effective dead-time management. We propose to set the level of 'acceptable dead time' to less than 10% at 1 kHz trigger rate, because we should assume that all systems in DAQ1000 should be laid out for a 1 kHz trigger rate (even if this might not be fully achievable for the TPC). Simulations in the proposal show that this acceptable dead time can be accomplished with the proposed readout chain.

e.) There was some ambiguity regarding the actual proposed layout of the RDO chain. The committee determined that each SSD ladder will be read out by two ADC boards, one on each side of a ladder. This means that the signal will be routed to both sides of the TPC, which also means that cables and cooling lines will have to be mounted on the old SVT cone on the East side, and the new FGT support cone on the West side. The ADC boards on each side will be routed to four (4) RDO boards which are located in a single RDO crate residing outside the pole tips. So there will be two RDO crates, one on the East and one on the West side, each holding a total of four RDO boards. Each RDO board will have its own DDL link. The committee expressed concern regarding the long-term availability of PCI-X motherboards for handling the eight DDL links, simply because PCI-X is not considered the latest industry standard anymore (PCI-Express is the new standard). It is probable that the DDL receiver board will become available in PCI Express but not likely in time for 2010 installation. The DDL processing boards (D-RORC) and the SIU which resides at the detector are now industry standards and commercially available through CERNTECH in Budapest. We urge the SSD group to inquire on the potential upgrade of the boards to PCI-E, if necessary. We also note that at 1 kHz trigger rate, even the full raw data volume loads the eight DDL links to 41% of their capacity. It should be considered to use only four DDL links (one per pair of RDO's), especially if zero-suppression will be implemented on the RDO. There should be no impact on dead time, assuming the RDO buffers the data as recommended.

f.) Regarding the radiation length budget, the re-location of the RDO crates to behind the pole tips solves some of the issues in forward direction. A detailed off-detector budget, in particular along the support cones is not available, though. It needs to be produced as

soon as the readout chain and cooling system design issues have been resolved, and then it needs to be reviewed by the EEMC, FGT, and FMS detector groups. We don't expect any showstoppers in terms of radiation length but the forward detector groups should be made aware of the additional radiation length budget.

g.) Regarding the upgrade of the TCD in STAR we recommend that the SSD group designs its system assuming the current TCD architecture. Any major TCD upgrade, which might include hardware and protocol changes, needs to be implemented in all STAR systems. A TCU/TCD design meeting is scheduled for April 2008 at BNL, which should yield more specific information, but any new TCD will be required to support any STAR legacy equipment and therefore need not affect the SSD upgrade.

h.) The detector did not perform well in Run-7 (around 50% live). Many of the problems have likely occurred because of inhibited cooling, but the group should assess the quality of the existing ladders in tests at Nantes prior to a technical review later this year.

Detector Integration

The impact of integrating the new readout and cooling into STAR can only be fully assessed after the RDO chain and cooling system has been fully designed. There are no obvious show-stoppers, but a technical review committee needs to thoroughly review the integration work, as soon as the design has been completed. The committee strongly recommends a close collaboration with the HFT/FGT engineers who are charged with developing the new West side support cone in STAR. Cable and cooling routing as well as mounting structures need to be integrated, ideally by a dedicated integration engineer which should be provided by the HFT group. Final integration plans also require review by the forward detector groups (FGT, FMS, EEMC). A potential common cooling system integration scheme between SSD and the IST part of the HFT needs to be explored and defined.
The DAQ and TRG experts on our committee have reviewed the integration issues regarding DAQ and TRG. We have not found any show-stoppers outside the current inability of the RDO firmware to handle aborts properly. A more detailed technical review is needed when the final RDO design is completed, with particular attention focused on the functionality of the TRG/SSD interface.

One major integration decision which resulted from the review deliberations is that the SSD group has agreed to place the RDO boxes in the new SVT-RDO location, i.e. outside the pole tips, for better maintenance during the run. The group also anticipates using the existing SVT water cooling for the RDO boxes, which is a sensible proposal.

Resources, Cost, and Schedule

The cost estimate seems reasonable based on past experiences by the committee members. But the final cost of such a project can only be assessed as soon as the final

design is developed. More effort could have been put into the cost estimate simply by deciding on a final, commercially available, ADC. If the power consumption would require any cooling other than air cooling, the cost of the project will increase considerably. We therefore propose to have a cost and schedule review in conjunction with the final technical review by the end of this year.

The resource loading is actually well established. Additional information provided by the group, in particular with respect to personnel at Nantes helped in convincing the committee that the proposed resources are sufficient.

According to the SSD group the manpower available at Nantes is the following: Stephane Bouvier: project engineer : 10% (ladder repair, supervising the upgrade and testing the ladders), Gerard Guilloux: mechanical engineer : 5% (mounting/dismounting SSD and ladders), Christophe Renard : electronic engineer : less than 50% (upgrade designing), and Louis-Marie Rigalleau: electronic technician: 30% (ladder repair and testing). These resources at Nantes are essential and we strongly urge STAR to enter into a provisional MoU with Nantes prior to funding the project. The resources at BNL and LBNL seem to be sufficient as well, although the project would benefit from actually identifying the BNL engineer and LBNL resources who will be responsible for board layout, routing and PCB fabrication, prior to funding. We also recommend that STAR operations considers an increase in its FTE loading by about 0.1 FTE per year due to the SSD operation in the out years.

One of the biggest concerns of the committee is the complete lack of a management structure for the project. It is our recommendation that the project should only go forward after a well defined structure has been developed including not only a project leader, but also a chief engineer, a designer, and a liaison engineer to the HFT and FGT projects. Certainly for a project of such a limited scope, the engineering tasks could be taken over by a single engineer, and also the integration tasks could be handled by a HFT or FGT engineer, but all of these responsibilities have to not only be defined but also specific names need to be assigned in order to allow us to review the management structure in detail, Without it the lines of responsibilities are not defined at all, neither within the project nor with respect to STAR. We anticipate that scientists at LBNL will take on a major management role in the SSD, but we need to review the structure further after the group has agreed on it.

The technically driven schedule seems feasible and it is also necessary in order to keep the main task of the SSD, as a support detector for HFT and FGT, on track. Therefore we propose that the project receives its first batch of R&D funding this year and then undergoes a detailed technical, integration and management review by the end of the year. The rather rudimentary proposal, which has been reviewed here, should be updated by a detailed technical and management (incl. cost and schedule) proposal prior to this second review.